arXiv:1511.05266v1 [cs.LG] 17 Nov 2015

Semi-supervised Collaborative Ranking with Push at Top

Iman Barjasteh[†*], Rana Forsati[‡*], Abdol-Hossein Esfahanian[‡], Hayder Radha[†]

[†]Department of Electrical and Computer Engineering

[‡]Department of Computer Science and Engineering

Michigan State University

{forsati,esfahanian}@cse.msu.edu, {barjaste,radha}@msu.edu

Abstract

Existing collaborative ranking based recommender systems tend to perform best when there is enough observed ratings for each user and the observation is made completely at random. Under this setting recommender systems can properly suggest a list of recommendations according to the user interests. However, when the observed ratings are extremely sparse (e.g. in the case of cold-start users where no rating data is available), and are not sampled uniformly at random, existing ranking methods fail to effectively leverage side information to transduct the knowledge from existing ratings to unobserved ones. We propose a **semi-supervised collaborative ranking** model, dubbed S^{2}COR, to improve the quality of cold-start recommendation. S^{2}COR mitigates the sparsity issue by leveraging side information about *both observed and missing* ratings by collaboratively learning the ranking model. This enables it to deal with the case of missing data not at random, but to also effectively incorporate the available side information in transduction. We experimentally evaluated our proposed algorithm on a number of challenging real-world datasets and compared against state-of-the-art models for cold-start recommendation. We report significantly higher quality recommendations with our algorithm compared to the state-of-the-art.

1 Introduction

Due to the popularity and exponential growth of e-commerce and online streaming websites, a compelling demand has been created for efficient recommender systems to guide users

*These authors contributed equally to this work.

toward items of their interests (e.g. products, books, movies) [1]. In collaborative methods, either *filtering* or *ranking*, by relying on the low-rank assumption on the users' preferences, both users and items are mapped into a latent feature space based on partially observed ratings that are later used to make predictions. In collaborative filtering (CF) methods such as matrix factorization [15], where the aim is to accurately predict the ratings, the latent features are extracted in a way to minimize the prediction error measured in terms of popular performance measures such as root mean square error (RMSE). In spark contrast to CF, in collaborating ranking (CR) models [15, 10, 36, 11], where the goal is to rank the unrated items in the order of relevance to the user, the popular ranking measures such as as discounted cumulative gain (DCG), normalized discounted cumulative gain (NDCG), and average precision (AP) [14] are often employed to collaboratively learn a ranking model for the latent features.

Recent studies have demonstrated that CR models lead to significantly higher ranking accuracy over their traditional CF counterparts that optimize rating prediction. This is important considering the fact that what we really care in recommendation is not the actual values of ratings, but the order of items to be recommended to a specific user. Therefore, the error measures such as RMSE are often hopelessly insufficient, as their place equal emphasis on all the ratings. Among ranking models, the methods that mainly concentrate on the *top of the list* have received a considerable amount of attention, due to the higher probability of examining the top portion of the list of recommendations by users. Therefore, the introduction of ranking metrics such as push norm or infinite norm [26, 3, 10, 16], sparked a widespread interest in CR models and has been proven to be more effective in practice [35, 10].

Although CR models for recommender systems has been studied extensively and some progress has been made, however, the state of affairs remains unsettled: the issue of handling *cold-start* items in ranking models and coping with *not missing at random* assumption of ratings are elusive open issues. First, in many real world applications, the rating data are very sparse (e.g., the density of the data is around 1% for many publicly available datasets) or for a subset of users or items the rating data is entirely missing (knows as cold-start user and cold-start item problem, respectively) [28]. Second, collaborative filtering and ranking models rely on the critical assumption that the missing ratings are sampled uniformly at random. However, in many real applications of recommender systems, this assumption is not believed to hold, as invariably some users are more active than others and some items are rated by many people while others are rarely rated [34]. These issues have been investigated in factorization based methods, nonetheless, it is not straightforward to adapt them to CR models and are left open [10]. Motivated by these challenges, we ask the following fundamental question in the context of collaborative ranking models:

> *Is it possible to effectively learn a collaborative ranking model in the presence of cold-start items/users that is robust to the sampling of observed ratings?*

In this paper, we give an affirmative answer to the above question. In particular, we introduce a *semi-supervised collaborative ranking* model, dubbed $\mathtt{S^2COR}$, by leveraging side information about *both observed and missing ratings* in collaboratively learning the ranking

model. In the learned model, unrated items are conservatively pushed after the relevant and before the irrelevant items in the ranked list of items for each individual user. This crucial difference greatly boosts the performance and limits the bias caused by learning only from sparse non-random observed ratings. We also introduce a graph regularization method to exploit the side information about users to overcome the cold-start users problem. In summary, the key features of $\texttt{S}^2\texttt{COR}$ are:

- Inspired by recent developments in ranking at top [26, 3, 17], the proposed model is a collaborative ranking model that primarily focuses on the top of the recommendation list for each user. Moreover, in stark contrast to pairwise ranking models which have quadratic dependency on the number of items, the proposed ranking model has a linear dependency on the number of items, making it suitable for large-scale recommendation.
- It leverages side information about items with both observed and missing ratings while collaboratively learning the ranking model, which enables it to effectively incorporate the available side information in knowledge transduction.
- By incorporating the unrated items in ranking, it limits the bias caused by learning solely based on the observed ratings and consequently deals with the not missing at random issue of ratings.
- It is also able to leverage the similarity information between users based on a graph regularization method to make high quality recommendations for users with few ratings or cold-start users without an rating information.

To build the intuition on how incorporating missing ratings in $\texttt{S}^2\texttt{COR}$ is beneficial in handling cold-start problem and mitigating data sparsity issue, we note that in many real world applications the available feedback on items is extremely sparse, and therefore the ranking models fail to effectively leverage the available side information in transdcuting the knowledge from existing ratings to unobserved ones. This problem becomes especially eminent in cases where surrogate ranking models such as pairwise models are used due to their computational virtues, where the unobserved ratings do not play any role in learning the model. As a result, by leveraging rich sources of information about all items, one can potentially bridge the gap between existing items and new items to overcome the cold-start problem.

Turning to the non-random sampling issue of observed ratings, we note that the non-randomness is observing the ratings creates a bias in learning the model that negatively impacts the future predictions and may degrade the resulting recommendation accuracy if ignored. Therefore, the nature of missing ratings has to be modeled precisely as to obtain correct results. To reduce the effect of bias, the proposed ranking model takes a conservative approach and pushes the items with unknown ratings to the middle of ranked list, i.e., after the relevant and before the irrelevant items. This is equivalent to assuming a prior about the unknown ratings which is believed to perform well as investigated in [12]. However, unlike [12], the proposed ranking idea is free of deciding an explicit value for missing ratings

which makes it more valuable from a practical point of view.

We conduct thorough experiments on real datasets and compare our results with the state-of-the-art models for cold-start recommendation to demonstrate the effectiveness of our proposed algorithm in recommendation at the top of the list and mitigating the data sparsity issue. Our results indicate that our algorithm outperforms other algorithms and provides recommendation with higher quality compared to the state-of-the-art methods.

Organization. This paper is organized as follows. We briefly review related work in Section 2. We establish the notation and formally define the problem in Section 3. In Section 4, we propose the semi-supervised collaborative ranking model with a push at the top of the list. Section 5 discusses efficient convex and non-convex optimization algorithms for optimization. Section 6 generalizes the proposed algorithm to leverage similarity information about users. We empirically evaluate the proposed method in Section 7, and conclude in Section 8.

2 Related Work

There is now a vast body of literature on ranking models for recommendation, coping with non-random missing ratings and an even bigger body of literature on handling cold-start problem by leveraging side information; we restrict our literature review here to papers that are most directly related.

Collaborative ranking for recommendation. The last few years have seen a resurgence in collaborative ranking centered around the technique of exploiting low-rank structures, an approach we take as well. Several approaches to CR have recently been proposed that are mainly inspired by the analogy between query-document relations in IR and user-item relations in recommender systems. The PMF-based approach [4] uses the latent representations produced by matrix factorization as user-item features and learns a ranking model on these features. CofiRank [37] learns latent representations that minimize a ranking-based loss instead of the squared error. ListRankMF [30] aims at minimizing the cross entropy between the predict item permutation probability and true item permutation probability. In [16] a method for Local Collaborative Ranking (LCR) where ideas of local low-rank matrix approximation were applied to the pairwise ranking loss minimization framework is introduced. In [35] a framework that allows for pointwise as well as listwise training with respect to various ranking metrics is proposed. Finally, [10] proposed a CR model build on the recent developments in ranking methods [3, 26] that focus on accuracy at top and proposed CR methods with p-push and infinite push norms. Incorporating side information in their model which is left as open issue was the main motivation of the current work.

Cold-start recommendation with side information. Due in part to its importance, there has been an active line of work to address difficulties associated with cold-start users and items, where a common theme among them is to exploit auxiliary information about

users or items besides the rating data that are usually available [31]. A feature based regression ranking model for predicting the values (rates) of user-item matrix in cold-start scenarios by leveraging all information available for users and items is proposed in [24]. The kernelized matrix factorization approach studied in [38], which incorporates the auxiliary information into the MF. In [27] joint factorization of the user-item and item-feature matrices by using the same item latent feature matrix in both decompositions is utilized. The FBSM model is introduced in [29], which learns factorized bilinear similarity model for new items, given the rating information as well as the features of these items. Recently, [5] proposed a decoupling approach to transduct knowledge from side information to rating prediction which is able to handle both cold-start items and users problems in factorization based models.

Recommendation with not missing at random ratings. Substantial evidence for violations of the missing at random condition in recommender systems is reported in [20] and it has been showed that incorporating an explicit model of the missing data mechanism can lead to significant improvements in prediction performance.The first study of the effect of non-random missing data on collaborative ranking is presented in [19]. In [32] an EM algorithm to optimize in turn the factorization and the estimation of missing values. Recently, in [12] a novel dynamic matrix factorization framework that allows to set an explicit prior on unknown values is introduced. However their algorithm requires a careful setting of the prior rating to be practical.

3 Preliminaries

In this section we establish the notation used throughout the paper and formally describe our problem setting.

Scalars are denoted by lower case letters and vectors by bold face lower case letters such as $\mathbf{u}$. We use bold face upper case letters such as $\mathbf{M}$ to denote matrices. The Frobenius norm of a matrix $\mathbf{M} \in \mathbb{R}^{n\times m}$ is denoted by $\|\mathbf{M}\|_{\mathrm{F}}$, i.e, $\|\mathbf{M}\|_{\mathrm{F}} = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{m}|M_{ij}|^2}$ and its (i,j)th entry is denoted by $A_{i,j}$. The trace norm of a matrix is denoted by $\|\mathbf{M}\|_*$ which is defined as the sum of its singular values. The transpose of a vector and a matrix denoted by $\mathbf{u}^\top$ and $\mathbf{U}^\top$, respectively. We use $[n]$ to denote the set on integers $\{1, 2, \cdots, n\}$. The set of non-negative real numbers is denoted by $\mathbb{R}_+$. The indicator function is denoted by $\mathbb{I}[\cdot]$. For a vector $\mathbf{u} \in \mathbb{R}^p$ we use $\|\mathbf{u}\|_1 = \sum_{i=1}^{p}|\mathbf{u}_i|$, $\|\mathbf{u}\|_2 = \left(\sum_{i=1}^{p}|\mathbf{u}_i|^2\right)^{1/2}$, and $\|\mathbf{u}\|_\infty = \max_{1\le i\le p}\mathbf{u}_i$ to denote its ℓ_1, ℓ_2, and ℓ_∞ norms, respectively. The dot product between two vectors $\mathbf{u}$ and $\mathbf{u}'$ is denoted by either $\langle \mathbf{u}, \mathbf{u}'\rangle$ or $\mathbf{u}^\top\mathbf{u}'$.

In collaborative filtering we assume that there is a set of n users $\mathcal{U} = \{u_1, \cdots, u_n\}$ and a set of m items $\mathcal{I} = \{i_1, \cdots, i_m\}$ where each user u_i expresses opinions about a set of items. The rating information is summarized in an $n \times m$ matrix $\mathbf{R} \in \{-1, +1, ?\}^{n\times m}, 1 \le i \le n, 1 \le j \le m$ where the rows correspond to the users and the columns correspond to the items and (p, q)th entry is the rate given by user u_p to the item i_q. We note that the rating

matrix is partially observed and it is sparse in most cases. We are mainly interested in recommending a set of items for an active user such that the user has not rated these items before.

4 Transductive Collaborating Ranking

We now turn our attention to the main thrust of the paper where we present our transductive collaborative ranking algorithm with accuracy at top by exploiting the features of unrated data. We begin with the basic formulation and then extend it to incorporate the unrated items. The pseudo-code of the resulting learning algorithm is provided in Algorithm 1.

4.1 A basic formulation

We consider a ranking problem, where, given a set of users $\mathcal{U}$ and known user feedback on a set of items $\mathcal{I}$, the goal is to generate rankings of unobserved items, adapted to each of the users' preferences. Here we consider the bipartite setting in which items are either relevant (positive) or irrelevant (negative). Many ranking methods have been developed for bipartite ranking, and most of them are essentially based on pairwise ranking. These algorithms reduce the ranking problem into a binary classification problem by treating each relevant/irrelevant instance pair as a single object to be classified [18].

As mentioned above, most research has concentrated on the rating prediction problem in CF where the aim is to accurately predict the ratings for the unrated items for each user. However, most applications that use CF typically aim to recommend only a small ranked set of items to each user. Thus rather than concentrating on rating prediction we instead approach this problem from the ranking viewpoint where the goal is to rank the unrated items in the order of relevance to the user. Moreover, it is desirable to concentrate aggressively on top portion of the ranked list to include mostly relevant items and push irrelevant items down from the top. Specifically, we propose an algorithm that maximizes the number of relevant items which are pushed to the absolute top of the list by utilizing the P-Norm Push ranking measure which is specially designed for this purpose [26] .

For simplicity of exposition, let us first consider the ranking model for a single user u. Let $\mathcal{X}^+ = \{\mathbf{x}_1^+, \cdots, \mathbf{x}_{n_+}^+\}$ and $\mathcal{X}^- = \{\mathbf{x}_1^-, \cdots, \mathbf{x}_{n_-}^-\}$ be the set of feature vectors of n_+ relevant and n_- irrelevant items to user u, respectively. We consider linear ranking functions where each item features vector $\mathbf{x} \in \mathbb{R}^d$ is mapped to a score $\mathbf{w}^\top \mathbf{x}$. The goal is to find parameters $\mathbf{w}$ for each user such that the ranking function best captures past feedback from the user. The goal of ranking is to maximize the number of relevant items ranked above the highest-ranking irrelevant item. We cast this idea for each user u individually into the following optimization problem:

$$\min_{\mathbf{w}\in\mathbb{R}^d} \frac{1}{n^+} \sum_{i=1}^{n^+} \mathbb{I}\left[\langle \mathbf{w}, \mathbf{x}_i^+ \rangle \leq \max_{1\leq j\leq n^-} \langle \mathbf{w}, \mathbf{x}_j^- \rangle\right] \tag{1}$$

Algorithm 1 S^{2}COR

input: $\lambda \in \mathbb{R}_+$: the regularization parameter, and $\{\eta_t\}_{t\geq 1}$: the sequence of scalar step sizes
Initialize $\mathbf{W}_0 \in \mathbb{R}^{n\times d}$
Choose an appropriate step size
for $t = 1, \ldots, T$ **do**
  Compute the sub-gradient of $\mathbf{G}_t \in \partial\mathcal{L}(\mathbf{W}_t)$ using Eq. (11)
  $[\mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t] \leftarrow \texttt{SVD}(\mathbf{W}_{t-1} - \frac{1}{\eta_{t-1}}\mathbf{G}_t))$
  $\mathbf{W}_t \leftarrow \mathbf{U}_t \left[\mathbf{\Sigma} - \frac{\lambda}{\eta_{t-1}}\mathbf{I}\right]_+ \mathbf{V}_t^\top$
end for
output:

where $\mathbb{I}[\cdot]$ is the indicator function which returns 1 when the input is true and 0 otherwise, n^+ and n^- are the the number of relevant and irrelevant items to user u, respectively.

Let us now derive the general form of our objective. We hypothesize that most users base their decisions about items based on a number of latent features about the items. In order to uncover these latent feature dimensions, we impose a low-rank constraint on the set of parameters for all users. To this end, let $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \cdots, \mathbf{w}_n]^\top \in \mathbb{R}^{n\times d}$ denote the matrix of all parameter vectors for n users. Let $\mathcal{I}_i^+ \subseteq \{1, 2, \ldots, m\}$ and $\mathcal{I}_i^- \subseteq \{1, 2, \ldots, m\}$ be the set of relevant and irrelevant items of ith user, respectively. The overall objective for all users is formulated as follows:

$$\mathcal{F}(\mathbf{W}) = \lambda\|\mathbf{W}\|_* + \sum_{i=1}^{n}\left(\frac{1}{|\mathcal{I}_i^+|}\sum_{j\in\mathcal{I}_i^+}\mathbb{I}\left[\langle\mathbf{w}_i, \mathbf{x}_j\rangle \leq \max_{k\in\mathcal{I}_i^-}\langle\mathbf{w}_i, \mathbf{x}_k\rangle\right]\right), \tag{2}$$

where $\|\cdot\|_*$ is the trace norm (also known as nuclear norm) which is the sum of the singular values of the input matrix.

The objective in Eq. (2) is composed of two terms. The first term is the regularization term and is introduced to capture the factor model intuition discussed above. The premise behind a factor model is that there is only a small number of factors influencing the preferences, and that a user's preference vector is determined by how each factor applies to that user. Therefore, the parameter vectors of all users must lie in a low-dimensional subspace. Trace-norm regularization is a widely-used and successful approach for collaborative filtering and matrix completion. The trace-norm regularization is well-known to be a convex surrogate to the matrix rank, and has repeatedly shown good performance in practice [33, 9]. The second term is introduced to push the relevant items of each user to the top of the list when ranked based on the parameter vector of the user and features of items.

The above optimization problem is intractable due to the non-convex indicator function. To design practical learning algorithms, we replace the indicator function in (2) with its convex surrogate. To this end, define the convex loss function $\ell : \mathbb{R} \mapsto \mathbb{R}_+$ as $\ell(x) =$

$[1-x]_+$. This is the widely used hinge loss in SVM classification (see e.g., [7]) [1]. This loss function reflects the amount by which the constraints are not satisfied. By replacing the non-convex indicator function with this convex surrogate leads to the following tractable convex optimization problem:

$$\begin{aligned}\mathcal{F}(\mathbf{W}) &= \lambda\|\mathbf{W}\|_* \\ &+ \sum_{i=1}^{n}\left(\frac{1}{|\mathcal{I}_i^+|}\sum_{j\in\mathcal{I}_i^+}\ell\left(\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty\right)\right)\end{aligned} \tag{3}$$

where $\mathbf{X}_i^- = [\mathbf{x}_1,\ldots,\mathbf{x}_{n_i^-}]^\top$ is the matrix of features of n_i^- irrelevant items in $\mathcal{I}_i^-$ and $\|\cdot\|_\infty$ is the max norm of a vector.

4.2 Semi-supervised collaborative ranking

In this part, we extend the proposed ranking idea to learn both from rated as well as unrated items. The motivation of incorporating unrated items comes from the following key observations. First, we note that commonly there is a small set of rated (either relevant or irrelevant) items for each user and a large number of unrated items. As it can be seen from Eq. (2), the unrated items do not play any role in learning the model for each user as the learning is only based on the pair of rated items. When the feature information for items is available, it would be very helpful if one can leverage such unrated items in the learning-to-rank process to effectively leverage the available side information. By leveraging both types of rated and unrated items, we can compensate for the lack of rating data. Second, the non-randomness in observing the observed ratings creates a bias in learning the model that may degrade the resulting recommendation accuracy. Therefore, finding a precise model to reduce the effect of bias introduced by non-random missing ratings seems essential.

To address these two issues, we extend the basic formulation in Eq. (2) to incorporate items with missing ratings in ranking of items for individual users. A conservative solution is to push the items with unknown ratings to the middle of ranked list, i.e., after the relevant and before the irrelevant items. To do so, let $\mathcal{I}_i^\circ = \mathcal{I}\setminus\left(\mathcal{I}_i^+\cup\mathcal{I}_i^-\right)$ denote the set of items unrated for user $i\in\mathcal{U}$. We introduce two extra terms in the objective in Eq. (2) to push the unrated items $\mathcal{I}_\circ^i$ below the relevant items and above the irrelevant items, which yilelds

[1] We note that other convex loss functions such as exponential loss $\ell(x)=\exp(-x)$, and logistic loss $\ell(x)=\log(1+\exp(-x))$ also can be used as the surrogates of indicator function, but for the simplicity of derivation we only consider the hinge loss here.

the following objective:

$$\begin{aligned}\mathcal{L}(\mathbf{w}) = \quad & \frac{1}{|\mathcal{I}_i^+|}\sum_{i\in\mathcal{I}_i^+}\ell\left(\langle\mathbf{w},\mathbf{x}_i\rangle \le \max_{j\in\mathcal{I}_i^-}\langle\mathbf{w},\mathbf{x}_j\rangle\right) \\ + \quad & \frac{1}{|\mathcal{I}_i^+|}\sum_{i\in\mathcal{I}_i^+}\ell\left(\langle\mathbf{w},\mathbf{x}_i\rangle \le \max_{j\in\mathcal{I}_i^\circ}\langle\mathbf{w},\mathbf{x}_j\rangle\right) \\ + \quad & \frac{1}{|\mathcal{I}_i^\circ|}\sum_{i\in\mathcal{I}_i^\circ}\ell\left(\langle\mathbf{w},\mathbf{x}_i\rangle \le \max_{j\in\mathcal{I}_i^-}\langle\mathbf{w},\mathbf{x}_j\rangle\right)\end{aligned} \tag{4}$$

Equipped with the objective of individual users, we now turn to the final collaborating ranking objective as:

$$\begin{aligned}\mathcal{F}(\mathbf{W}) &= \lambda\|\mathbf{W}\|_* \\ &+ \sum_{i=1}^n\left(\frac{1}{|\mathcal{I}_i^+|}\sum_{j\in\mathcal{I}_i^+}\ell\left(\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty\right)\right) \\ &+ \sum_{i=1}^n\left(\frac{1}{|\mathcal{I}_i^+|}\sum_{j\in\mathcal{I}_i^+}\ell\left(\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^\circ\mathbf{w}_i\|_\infty\right)\right) \\ &+ \sum_{i=1}^n\left(\frac{1}{|\mathcal{I}_i^\circ|}\sum_{j\in\mathcal{I}_i^\circ}\ell\left(\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty\right)\right),\end{aligned} \tag{5}$$

where $\mathbf{X}_i^\circ = [\mathbf{x}_1,\ldots,\mathbf{x}_{n_i^\circ}]^\top$ is the matrix of n_i° unrated items in $\mathcal{I}_i^\circ$.

Remark 4.1 *We emphasize that beyond the accuracy considerations of push norm at top of the list, the push norm has a clear advantage to the pairwise ranking models from a computational point of view. In particular, the push norm has a linear $O(m)$ dependency on the number of items which is quadratic $O(m^2)$ for pairwise ranking models.*

5 The Optimization

We now turn to solving the optimization problem in (5). We start by discussing a gradient descent method with shrinkage operator followed by its accelerated version, and then propose a non-convex formulation with alternative minimization for more effective optimization of objective in S^{2}COR.

5.1 Gradient descent with shrinkage operator

Due to the presence of trace norm of the parameters matrix, this objective function falls into the general category of composite optimization, which can be solved by stochastic gradient

or gradient descent methods. In this part we propose a projected gradient decent method to solve the optimization problem. First we write the objective as:

$$\min_{\mathbf{W}\in\mathbb{R}^{n\times d}} \mathcal{F}(\mathbf{W}) = \lambda\|\mathbf{W}\|_* + \mathcal{L}(\mathbf{W}), \tag{6}$$

where $\mathcal{L}(\mathbf{W}) = \sum_{i=1}^{n} \mathcal{L}(\mathbf{w}_i)$.

A simple way to solving the above optimization problem is gradient descent algorithm [23], which needs to evaluate the gradient of objective at each iteration. To deal with the non-smooth trace norm $\|\mathbf{W}\|_*$ in the objective, we first note that the optimization problem in Eq. (6) can be reformulated under the framework of proximal regularization or composite gradient mapping [23]. By taking advantage of the composite structure it is possible to retain the same convergence rates of the gradient method for the smooth optimization problems. In particular, the optimization problem in (6) can be solved iteratively by:

$$\begin{aligned}\mathbf{W}_t = \arg\min_{\mathbf{W}} \mathcal{L}(\mathbf{W}_{t-1}) + \mathtt{tr}\left((\mathbf{W}-\mathbf{W}_{t-1})^\top \nabla\mathcal{L}(\mathbf{W}_{t-1})\right) \\ + \frac{\eta_t}{2}\|\mathbf{W}-\mathbf{W}_{t-1}\|_{\mathrm{F}}^2 + \lambda\|\mathbf{W}\|_*,\end{aligned} \tag{7}$$

where $\{\eta_t\}_{t\geq 1}$ is a sequence of scalar step sizes and $\mathtt{tr}(\cdot)$ is the trace of input matrix.

By ignoring the constant terms, Eq. (7) can also be rewritten as:

$$\frac{\eta_t}{2}\left\|\mathbf{W} - \left(\mathbf{W}_{t-1} - \frac{1}{\eta_t}\nabla\mathcal{L}(\mathbf{W}_{t-1})\right)\right\|_{\mathrm{F}}^2 + \lambda\|\mathbf{W}\|_*. \tag{8}$$

We use the singular value shrinkage operator introduced in [8] to find the optimal solution to Eq. (8). To this end, consider the singular value decomposition (SVD) of a matrix $\mathbf{M}\in\mathbb{R}^{n\times d}$ of rank r as $\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, $\quad \mathbf{\Sigma} = \mathrm{diag}(\{\sigma_i\}_{1\leq i\leq r})$, where $\mathbf{U}$ and $\mathbf{V}$ are respectively $n\times r$ and $d\times r$ matrices with orthonormal columns, and the singular values σ_i are positive. For a scalar $\tau\in\mathbb{R}_+$, define the singular value shrinkage operator $\mathscr{P}_\tau$ as:

$$\mathscr{P}_\tau(\mathbf{M}) := \mathbf{U}\mathscr{P}_\tau(\mathbf{\Sigma})\mathbf{V}^*, \quad \mathscr{P}_\tau(\mathbf{\Sigma}) = \mathrm{diag}([\sigma_i - \tau]_+\}), \tag{9}$$

where $[x]_+$ is the positive part of x, namely, $[x]_+ = \max(0, x)$. The shrinkage operator basically applies a soft-thresholding rule to the singular values of $\mathbf{M}$, effectively shrinking these towards zero.

Theorem 5.1 (Theorem 2.1, [8]) *For each $\tau \geq 0$ and $\mathbf{W}\in\mathbb{R}^{n\times d}$, the singular value shrinkage operator* (9) *obeys*

$$\mathscr{P}_\tau(\mathbf{W}) = \arg\min_{\mathbf{X}} \left\{\frac{1}{2}\|\mathbf{X}-\mathbf{W}\|_{\mathrm{F}}^2 + \tau\|\mathbf{X}\|_*\right\}. \tag{10}$$

The above theorem shows that the singular value shrinkage operator of a matrix is the solution to Eq. (10). Equipped with this result, the optimization problem in Eq. (8) can be

solved by first computing the SVD of updated $\mathbf{W}_{t-1}$ and then applying soft thresholding on the singular values as:

$$\mathbf{W}_t = \mathscr{P}_{\frac{\lambda}{\eta_{t-1}}}\left(\mathbf{W}_{t-1} - \frac{1}{\eta_{t-1}}\nabla\mathcal{L}(\mathbf{W}_{t-1})\right)$$

Now we need to evaluate the gradient of $\mathcal{L}(\mathbf{W})$ at $\mathbf{W}_{t-1}$. The convex function $\mathcal{L}(\mathbf{W})$ is not differentiable, so we use its subgradient in updating the solutions which can be computed for ith parameter vector $\mathbf{w}_i$, as follows:

$$\begin{aligned}
\mathbf{g}_i &= \partial\mathcal{L}/\partial\mathbf{w}_i \\
&= \sum_{j\in\mathcal{I}_i^+} \mathbb{I}\left[\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty \le 1\right]\left(\partial\|\mathbf{X}_i^-\mathbf{w}_i\|_\infty - \mathbf{x}_j\right) \\
&+ \sum_{j\in\mathcal{I}_i^+} \mathbb{I}\left[\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^\circ\mathbf{w}_i\|_\infty \le 1\right]\left(\partial\|\mathbf{X}_i^\circ\mathbf{w}_i\|_\infty - \mathbf{x}_j\right) \\
&+ \sum_{j\in\mathcal{I}_i^\circ} \mathbb{I}\left[\langle\mathbf{w}_i,\mathbf{x}_j\rangle - \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty \le 1\right]\left(\partial\|\mathbf{X}_i^-\mathbf{w}_i\|_\infty - \mathbf{x}_j\right),
\end{aligned} \tag{11}$$

where $\partial\|\mathbf{X}_i^-\mathbf{w}_i\|_\infty$ is the subdifferential of the function $\|\mathbf{X}_i^-\mathbf{w}_i\|_\infty$ at point $\mathbf{w}_i$. Since the subdifferential of the maximum of functions is the convex hull of the union of subdifferentials of the active functions at point $\mathbf{w}$ [22], we have:

$$\begin{aligned}
\partial\|\mathbf{X}_i^-\mathbf{w}_i\|_\infty &= \partial \max_{1\le j\in\mathcal{I}_i^-}\langle\mathbf{w}_i,\mathbf{x}_j\rangle \\
&= \mathrm{conv}\left\{\mathbf{x}_j | \langle\mathbf{w}_i,\mathbf{x}_j\rangle = \|\mathbf{X}_i^-\mathbf{w}_i\|_\infty, j\in\mathcal{I}_j^-\right\}.
\end{aligned}$$

Then, $\mathbf{G} = [\mathbf{g}_1,\mathbf{g}_2,\ldots,\mathbf{g}_n]^\top \in \mathbb{R}^{n\times d}$ is a subgradient at $\mathbf{W}$, i.e. $\mathbf{G}\in\partial\mathcal{L}(\mathbf{W})$.

Remark 5.2 *We note that here, for the ease of exposition, we only consider the non-differentiable convex hinge loss as the surrogate of non-convex due to its computational virtues. However, by using other smooth convex surrogate losses such as smoothed hinge loss or exponential loss, one can apply the accelerated gradient descent methods [22] to solve the optimization problem which results in significantly faster convergence rate compared to the naive gradient descent (i.e, $O(1/\sqrt{\epsilon})$ convergence rate for accelerated method compared to the $O(1/\epsilon^2)$ rate for gradient descent for non-smooth optimization, where ϵ is the target accuracy).*

5.2 Efficient optimization by dropping convexity

The main computational cost in each iteration of the `S`^{2}`COR` algorithm lies in computing the SVD decomposition of $\mathbf{W}_k$. An alternative cheaper way to solving the optimization problem in Eq. (6) is as follows. For a fixed rank of the target parameter matrix $\mathbf{W}$, say k, one

can decompose it as $\mathbf{W} = \mathbf{UV}$. From the equivalence relation between trace norm and the Frobenius of its components in decomposition,

$$\|\mathbf{W}\|_* = \min_{\substack{\mathbf{U}\in\mathbb{R}^{n\times k},\mathbf{V}\in\mathbb{R}^{m\times k}\\ \mathbf{W}=\mathbf{U}\mathbf{V}^\top}} \frac{1}{2}\left(\|\mathbf{U}\|_{\mathrm{F}}^2 + \|\mathbf{V}\|_{\mathrm{F}}^2\right)$$

we can write the objective in terms of these low rank components as:

$$\min_{\mathbf{U}\in\mathbb{R}^{n\times k},\mathbf{V}\in\mathbb{R}^{m\times k}} \frac{\lambda}{2}\left(\|\mathbf{U}\|_{\mathrm{F}}^2 + \|\mathbf{V}\|_{\mathrm{F}}^2\right) + \mathcal{L}(\mathbf{UV}), \tag{12}$$

These factored optimization problem does not have the explicit trace norm regularization. However, the new formulation is non-convex and potentially subject to stationary points that are not globally optimal. However, despite its non-convexity, the formulation in Eq. (12) is competitive as compared to trace-norm minimization, while scalability is much better. In particular, the objective is not jointly convex in both $\mathbf{U}$ and $\mathbf{V}$ but it is convex in each of them fixing the other one. Therefore, to find a local solution one can stick to the standard gradient descent method to find a solution in an iterative manner as follows:

$$\mathbf{U}_{t+1} \leftarrow (1-\lambda\eta_t)\mathbf{U}_t - \eta_t \nabla_{\mathbf{U}}\mathcal{L}|_{\mathbf{U}=\mathbf{U}_t,\mathbf{V}=\mathbf{V}_t},$$
$$\mathbf{V}_{t+1} \leftarrow (1-\lambda\eta_t)\mathbf{V}_t - \eta_t \nabla_{\mathbf{V}}\mathcal{L}|_{\mathbf{U}=\mathbf{U}_t,\mathbf{V}=\mathbf{V}_t}.$$

Remark 5.3 *It is remarkable that the large number of users or items may cause computational problems in solving the optimization problem using GD method. The reason is essentially the fact that computing the gradient at each iteration requires to go through all the users and compute the gradient for pair of items. To alleviate this problem one can utilize stochastic gradient method [21] to solve the optimization problem. The main idea is to choose a fixed subset of pairs for gradient computation instead of all pairs at each iteration or a sample a user at random for gradient computation instead of including all users. We note that this strategy generates unbiased estimates of the true gradient and makes each iteration of algorithm computationally more efficient compared to the full gradient counterpart.*

6 Regularization by Exploiting Similarity of Users

In many recommender systems, in addition to handling cold-start items, it would be beneficial to make high quality recommendations to new users. The new user cold start issue represents a serious problem in recommender systems as it can lead to the loss of new users who decide to stop using the system due to the lack of accuracy in the recommendations received in that first stage in which they have not yet cast a significant number of ratings with which to feed the recommender system's collaborative filtering core.

The above formulation estimates a parameter vector separately for each user regardless of potential similarities across users. In this section we generalize the proposed ranking model

Algorithm 2 S^{2}COR+

input: $\gamma, \lambda \in \mathbb{R}_+$: the regularization parameters, $\{\eta_t\}_{t\geq 1}$: the sequence of scalar step sizes, and $\mathbf{S} \in \mathbb{R}^{n\times n}$: the similarity matrix of usets
Initialize $\mathbf{W}_0 \in \mathbb{R}^{n\times d}$
Choose an appropriate step size
for $t = 1, \ldots, T$ **do**
  Compute the sub-gradient of $\mathbf{G}_t \in \partial\mathcal{L}(\mathbf{W}_t)$ using Eq. (11)
  Compute $\widehat{\mathbf{G}}_t = \mathbf{G}_t + 2\gamma\mathbf{X}\mathbf{X}^\top\mathbf{W}\mathbf{L}$
  $[\mathbf{U}_t, \mathbf{\Sigma}_t, \mathbf{V}_t] \leftarrow \texttt{SVD}(\mathbf{W}_{t-1} - \frac{1}{\eta_{t-1}}\widehat{\mathbf{G}}_t))$
  $\mathbf{W}_t \leftarrow \mathbf{U}_t \left[\mathbf{\Sigma} - \frac{\lambda}{\eta_{t-1}}\mathbf{I}\right]_+ \mathbf{V}_t^\top$
end for
output:

to the setting where the similarity information between users is available. Let $\mathbf{S} \in \mathbb{R}^{n\times n}$ be the similarity matrix of users which is inferred from the side information about users such as social relations between them or explicit users' features. Let $\mathbf{D}$ be the diagonal matrix with $D_{ii} = \sum_{j=1}^{n} W_{ij}$, and $\mathbf{L} = \mathbf{D} - \mathbf{S}$ be the Laplacian matrix. It is natural to require the similar users in the matrix $\mathbf{S}$ have similarity evaluation on rating the items [6]. Thus, the new regularization on the parameters of users $\mathbf{W}$ using the similarity matrix $\mathbf{S}$ can be achieved by minimizing:

$$
\begin{aligned}
&\frac{1}{2}\sum_{i=1}^{m}\left(\sum_{j,k=1}^{n} S_{jk}\left(\langle \mathbf{w}_j, \mathbf{x}_i\rangle - \langle \mathbf{w}_k, \mathbf{x}_i\rangle\right)^2\right)\\
&= \frac{1}{2}\sum_{i=1}^{m}\left(\mathbf{x}_i^\top\mathbf{w}_j(\sum_{k} S_{jk})\mathbf{w}_j^\top\mathbf{x}_i - \sum_{j,k}\mathbf{x}_i^\top\mathbf{w}_j S_{jk}\mathbf{w}_k^\top\mathbf{x}_i\right)\\
&= \sum_{i}\left(\sum_{j}\mathbf{x}_i^\top\mathbf{w}_j S_{jj}\mathbf{w}_j^\top\mathbf{x}_i - \sum_{j,k}\mathbf{x}_i^\top\mathbf{w}_j S_{ij}\mathbf{w}_k^\top\mathbf{x}_i\right)\\
&= \sum_{i}\mathbf{x}_i^\top\left(\mathbf{W}\mathbf{D}\mathbf{W}^\top - \mathbf{W}\mathbf{S}\mathbf{W}^\top\right)\mathbf{x}_i\\
&= \sum_{i}\mathbf{x}_i^\top\mathbf{W}\mathbf{L}\mathbf{W}^\top\mathbf{x}_i = \texttt{tr}\left(\mathbf{X}^\top\mathbf{W}\mathbf{L}\mathbf{W}^\top\mathbf{X}\right).
\end{aligned}
$$

By plugging the above regularization term in the Eq. (6), we obtain the following optimization problem:

$$\mathcal{F}(\mathbf{W}) = \lambda\|\mathbf{W}\|_* + \mathcal{L}(\mathbf{W}) + \gamma\ \texttt{tr}\left(\mathbf{X}^\top\mathbf{W}\mathbf{L}\mathbf{W}^\top\mathbf{X}\right),$$

The above optimization problem can also be solved using the optimization procedure discussed before. The only difference is the gradient computation step which now has an extra term due the introduction of user's similarity regularization. Specifically, we have the new gradient $\widehat{\mathbf{G}}_t$ computed as:

$$\widehat{\mathbf{G}}_t = \mathbf{G}_t + 2\gamma \mathbf{X}\mathbf{X}^\top \mathbf{W}\mathbf{L},$$

which replaces the $\mathbf{G}_t$ in Algorithm 1. The resulting algorithm, dubbed $\mathtt{S}^2\mathtt{COR}+$, is detailed in Algorithm 2.

7 Experiments

In this section, we conduct exhaustive experiments to demonstrate the merits and advantages of the proposed algorithm. We conduct our experiments on three well-known datasets MovieLens, Amazon and CiteULike. We investigate how the proposed $\mathtt{S}^2\mathtt{COR}$ performs in comparison to the state-of-the-art methods. In the following, first we introduce the datasets that we use in our experiments and then the metrics that we employ to evaluate the results, followed by our detailed experimental results on the real datasets.

In the following subsections, we intend to answer these key questions:

- **Ranking versus rating:** How does learning optimization for a ranking based loss function affect the performance of recommending versus the square loss function?
- **Employing missing ratings:** How does employing the missing ratings could help in making more accurate recommendations?
- **Dealing with cold-start items:** How does the proposed algorithm, with incorporating side information of items, perform in comparison to the state-of-the-art algorithms to deal with cold-start items?

7.1 Datasets

We use the following well known datasets to evaluate the performance of $\mathtt{S}^2\mathtt{COR}$:

- **ML-IMDB.** We used ML-IMDB which is a dataset extracted from the IMDB and the MovieLens 1M datasets by mapping the MovieLens and IMDB and collecting the movies that have plots and keywords. The rating values are 10 discrete numbers ranging from 1 to 10 and the rating were made binary by treating all the ratings greater than 5 as $+1$ and below 5 as -1.

- **Amazon.** We used the dataset of best-selling books and their ratings in Amazon. Each book has a one or two paragraphs of textual description, which has been used to have a set of features of the books. Ratings can be integers numbers from 1 to 5. The ratings were also made binary by treating all the ratings greater or equal to 3 as $+1$ and below 3 as -1.

Table 1: Statistics of real datasets used in our experiments.

Statistics	ML-IMDB	Amazon	CiteULike
Number of users	2,113	13,097	3,272
Number of items	8,645	11,077	21,508
Number of ratings	739,973	175,612	180,622
Number of features	8,744	5,766	6,359
Average number of ratings by users	350.2	13.4	55.2
Average number of ratings for items	85.6	13.4	55.2
Density	4.05%	0.12%	0.13%

- **CiteULike.** It is an online free service for managing and discovering scholarly references. Users can add those articles that they are interested in to their libraries. Collected articles in a user's library will be considered as relevant items for that user. This dataset does not have explicit irrelevant items and was chosen to illustrate the effect of considering missing data while only having relevant itmes.

For all above datasets, the description about the items were tokenized and after removing the stop words, the rest of the words were stemmed. Then those words that have been appeared in less than 20 items and more that 20% of the items were also removed [29]. At the end, the TF-IDF was applied on the remaining words and the TF-IDF scores represented the features of the items. The statistics of the datasets are given in Table 1. As it is shown in Table 1, all these three datasets have high dimensional feature space.

7.2 Metrics

We adopt the widely used metrics, Discounted Cumulative Gain at n and Recall at n, for assessing the performance of our and baseline algorithms. For each user u, given an item i, let s_k be the relevance score of the item ranked at position k, where $s_k = 1$ if the item is relevant to the user u and $s_k = 0$ otherwise. Now, given the list of top-n item recommendations for user u, Discounted Cumulative Gain at n, is defined as:

$$\mathrm{DCG_u@}n = s_1 + \sum_{k=2}^{n} \frac{s_k}{\log_2(k)}$$

If we divide the $\mathrm{DCG_u@}n$ by its maximum value, we get the $\mathrm{NDCG_u@}n$ value. Given the list of top-n item recommendations for each user u, Recall at n will count the number of relevant items appeared in the recommendation list divided by the size of the list. Recall at n is defined as:

$$\mathrm{REC_u@}n = \frac{|\{\text{relevant items to } u\} \cap \{\text{top-}n \text{ recommended items}\}|}{|\{\text{top-}n \text{ recommended items}\}|}$$

DCG@n, $\mathrm{NDCG_u@}n$ and REC@n will be computed for each user and then will be averaged over all users.

7.3 Methodology

Given the partially observed rating matrix, we transformed the observed ratings of all datasets from a multi-level relevance scale to a two-level scale $(+1, -1)$ while 0 is considered for unobserved ratings. We randomly selected 60% of the observed ratings for training and 20% for validation set and consider the remaining 20% of the ratings as our test set. To better evaluate the results, we performed a 3-fold-cross validation and reported the average value for our results.

7.4 Baseline Algorithms

The proposed $\texttt{S}^2\texttt{COR}$ algorithm is compared to the following algorithms:

- **Matrix Factorization (MF) [33]:** Is a matrix completion method that factorizes the incomplete observed matrix and completes the matrix using the unveiled latent features.
- **Matrix Factorization with Side Information (KPMF) [38]:** Is a matrix completion based algorithm, which incorporates external side information of the users or items into the matrix factorization process. It imposes a Gaussian Process prior over all rows of the matrix, and the learned model explicitly captures the underlying correlation among the rows.
- **Decoupled Completion and Transduction (DCT) [5]:** Is a matrix factorization based algorithm that decouples the completion and transduction stages and exploits the similarity information among users and items to complete the (rating) matrix.
- **Feature Based Factorized Bilinear Similarity Model (FBS) [29]:** This algorithm uses bilinear model to capture pairwise dependencies between the features. In particular, this model accounts for the interactions between the different item features.
- **Collaborative User-specific Feature-based Similarity Models (CUFSM):** By using the history of ratings for users, it learns personalized user model across the dataset. This method is one of the best performing collaborative latent factor based model [13].
- **Regression based Latent Factor Model (RLF):**[2] This method incorporates the features of items in factorization process by transforming the features to the latent space using linear regression [2]. If the learning method is Markov Chain Monte Carlo, we name it RLF-MCMC.
- **Cosine Similarity Based Recommender (CSR):** Using the similarity between features of items, the preference score of a user on an item will be estimated.

We would like to mention that as baseline algorithms we only consider state-of-the art methods that are able to exploit the side information about items.

[2]The implementation of this method is available in LibFM library [25].

Table 2: Results of comparing rankings methods versus ratings methods on ML-IMDB. λ is regularization parameter, h is dimension of latent features, T is the number of iterations, η is learning rate and σ is the standard deviation.

Algorithms	**Hyperparameters**	**NDCG**@10
MF	$\eta = 0.003, h = 10, \sigma = 0.4$	0.0947
KPMF	$\eta = 0.003, h = 10, \sigma = 0.4$	0.1005
DCT	$h = 10$	0.0095
RLF-MCMC	$T = 100, \sigma = 0.1$	0.0248
$\mathtt{S^2COR}$	$\lambda = 0.55, h = 10, T = 100$	**0.265**

7.5 $\mathtt{S^2COR}$ vs. rating

Many different algorithms are trying to provide recommendations to users such that the predicted rating values be very close to the actual rates that users would provide. These algorithms try to minimize the error between the predicted values and actual rating values by minimizing Mean Squared Error (MAE) or Root Mean Square Error (RMSE) or etc. Then, due to the fact that users tend to only care about the top of their recommendation list, predicting a ranking of items of interest instead of ratings became the main focus of recent works [10]. In this section we compare the results of $\mathtt{S^2COR}$ with those state-of-the-art algorithms that try to predict ratings for unrated items. Among the state-of-the-art algorithms, we chose a diverse set of algorithms, which are Matrix Factorization, Matrix Factorization with Side Information, Decoupled Completion and Transduction and Regression Based Latent Factor Model. Table 2 shows the NDCG value of top 10 items of recommendation list. It shows that $\mathtt{S^2COR}$ outperformed all other rating prediction based algorithms in terms of NDCG measure. The results confirm the effectiveness of providing the ranks of items rather than their ratings.

7.6 Robustness to not missing at random ratings

In this section we compare the effect of incorporating the unobserved ratings in our learning in comparison with excluding them from our learning. Most of the methods in the literature ignore the unobserved ratings and train their model only base on observed ratings. By incorporating the unrated items in ranking, our method can limit the bias caused by learning solely based on the observed ratings and consequently deals with the not missing at random issue of ratings. Table 3 shows results of comparing these two scenarios for $\mathtt{S^2COR}$ on ML-IMDB. In order to see the difference between these two scenarios, we considered 70% of the ratings for training and 30% for test to have more ground truth for our testing. Table 3 shows the NDCG@5, 10,15 and 20 for both scenarios and it shows that incorporating the unobserved ratings causes to improve the accuracy of recommendation list. Hence, the NDCG values for top 5, 10, 15 and 20 items improved when unrated items were included as part of the training process.

Table 3: Results of employing missing ratings versus ignoring them on ML-IMDB. $\lambda = 0.6$ is regularization parameter, $h = 10$ is dimension of latent features, $T = 100$ is the number of iterations.

Algorithm: S^{2}COR	**NDCG@5**	**NDCG@10**	**NDCG@15**	**NDCG@20**
Observed ratings	1.1690	2.2218	2.8362	3.2849
Observed + missing ratings	**1.1794**	**2.2405**	**2.8585**	**3.3096**

7.7 Dealing with cold-start items

We now turn to evaluating the effectiveness of S^{2}COR for cold-start recommendation. To do so, we randomly selected 60% of the items as our training items and 20% for validation set and considered the remaining 20% of the items as our test set. In this scenario, baseline algorithms that are used for comparison are CSR, FBS, CUFSM and RLF. For the experiments, we used ML-IMDB, Amazon and CiteULike datasets. Table 4 shows the measurement results of applying mentioned algorithms on these datasets. For each test, the parameters' values producing the best ranking on the validation set were selected to be used and reported. As it can be seen from the results in Table 4, the proposed S^{2}COR algorithm outperformed all other baseline algorithms and provided a recommendations with higher quality in comparison to other methods. We can also see from the results of Table 4 that for the ML-IMDB dataset, the improvement in terms of REC@10 is significant compared to other datasets. Since the density of this dataset is much higher than other two datasets, this observation indicates that our method is more effective in utilizing side information compared to other methods. These results demonstrate the effectiveness of S^{2}COR in comparison with other state-of-the-art algorithms. S^{2}COR was able to outperform other state-of-the-art algorithms by considering the missing data and focusing on top of the recommendation list for cold-start items.

8 Conclusions

In this paper we introduced a semi-supervised collaborative ranking model by leveraging side information about both observed and missing ratings in collaboratively learning the ranking model. In the learned model, unrated items are conservatively pushed after the relevant and before the relevant items in the ranked list of items for each individual user. This crucial difference greatly boosts the performance and limits the bias caused by learning only from sparse non-random observed ratings. The proposed algorithm is compared with seven baseline algorithms on three real world datasets that demonstrated the effectiveness of proposed algorithm in addressing cold-start problem and mitigating the data sparsity problem, while being robust to sampling of missing ratings.

This work leaves few interesting directions as future work. First, we would like to investigate the performance of the proposed S^{2}COR algorithm when side information about

Table 4: Results on cold-start items. λ, μ_1 and β are regularization parameters, h is dimension of latent features, l is the number of similarity functions and T is the number of iterations.

	Algorithms	**Hyperparameters**	**DCG**@10	**REC**@10
ML-IMDB	CSR	—	0.1282	0.0525
	RLF	$h = 15$	0.0455	0.0155
	CUFSM	$l = 1, \mu_1 = 0.005$	0.2160	0.0937
	FBS	$\lambda = 0.01, \beta = 0.1, h = 5$	0.2270	0.0964
	$\texttt{S}^2\texttt{COR}$	$\lambda = 0.6, h = 10, T = 200$	**0.2731**	**0.2127**
Amazon	CSR	—	0.0228	0.1205
	RLF	$h = 30$	0.0076	0.0394
	CUFSM	$l = 1, \mu_1 = 0.25$	0.0282	0.1376
	FBS	$\lambda = 0.1, \beta = 1, h = 1$	0.0284	0.1392
	$\texttt{S}^2\texttt{COR}$	$\lambda = 0.6, h = 10, T = 200$	**0.1195**	**0.1683**
CiteULike	CSR	—	0.0684	0.1791
	RLF	$h = 75$	0.0424	0.0874
	CUFSM	$l = 1, \mu_1 = 0.25$	0.0791	0.2017
	FBS	$\lambda = 0.25, \beta = 10, h = 5$	0.0792	0.2026
	$\texttt{S}^2\texttt{COR}$	$\lambda = 0.6, h = 10, T = 200$	**0.0920**	**0.2243**

users is also available using the graph regularization idea discussed in Section 6. Second, we would like to empirically evaluate the performance of the optimization method derived by dropping the convexity in future. Also, we have largely ignored the case of differentiable smooth surrogate convex loss functions in this work and it would be interesting to consider smooth alternatives and apply the accelerated optimization methods for faster convergence. Moreover, the scalability analysis of of proposed algorithm on large datasets using stochastic optimization methods is also worthy of investigation. Finally, we believe there are still many open questions relating to non-random nature of rating information in many real applications.

References

[1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.

[2] D. Agarwal and B.-C. Chen. Regression-based latent factor models. In *SIGKDD*, pages 19–28. ACM, 2009.

[3] S. Agarwal. The infinite push: A new support vector ranking algorithm that directly optimizes accuracy at the absolute top of the list. In *SDM*, pages 839–850. SIAM, 2011.

[4] S. Balakrishnan and S. Chopra. Collaborative ranking. In *ACM WSDM*, pages 143–152. ACM, 2012.

[5] I. Barjasteh, R. Forsati, F. Masrour, A.-H. Esfahanian, and H. Radha. Cold-start item and user recommendation with decoupled completion and transduction. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 91–98. ACM, 2015.

[6] M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *The Journal of Machine Learning Research*, 7:2399–2434, 2006.

[7] C. J. Burges. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2):121–167, 1998.

[8] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

[9] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.

[10] K. Christakopoulou and A. Banerjee. Collaborative ranking with a push at the top. In *WWW*, pages 205–215. International World Wide Web Conferences Steering Committee, 2015.

[11] P. Cremonesi, Y. Koren, and R. Turrin. Performance of recommender algorithms on top-n recommendation tasks. In *ACM RecSys*, pages 39–46. ACM, 2010.

[12] R. Devooght, N. Kourtellis, and A. Mantrach. Dynamic matrix factorization with priors on unknown values. In *ACM SIGKDD*, pages 189–198. ACM, 2015.

[13] A. Elbadrawy and G. Karypis. User-specific feature-based similarity models for top-n recommendation of new items. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 6(3):33, 2015.

[14] K. Järvelin and J. Kekäläinen. Ir evaluation methods for retrieving highly relevant documents. In *ACM SIGIR*, pages 41–48. ACM, 2000.

[15] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.

[16] J. Lee, S. Bengio, S. Kim, G. Lebanon, and Y. Singer. Local collaborative ranking. In *Proceedings of the 23rd international conference on World wide web*, pages 85–96. ACM, 2014.

[17] N. Li, R. Jin, and Z.-H. Zhou. Top rank optimization in linear time. In *Advances in Neural Information Processing Systems*, pages 1502–1510, 2014.

[18] T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.

[19] B. M. Marlin and R. S. Zemel. Collaborative prediction and ranking with non-random missing data. In *RecSys*, pages 5–12. ACM, 2009.

[20] B. M. Marlin, R. S. Zemel, S. T. Roweis, and M. Slaney. Collaborative filtering and the missing at random assumption. In *UAI*, pages 267–275, 2007.

[21] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[22] Y. Nesterov. *Introductory lectures on convex optimization*, volume 87. Springer Science & Business Media, 2004.

[23] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.

[24] S.-T. Park and W. Chu. Pairwise preference regression for cold-start recommendation. In *RecSys*, pages 21–28. ACM, 2009.

[25] S. Rendle. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):57, 2012.

[26] C. Rudin. The p-norm push: A simple convex ranking algorithm that concentrates at the top of the list. *The Journal of Machine Learning Research*, 10:2233–2271, 2009.

[27] M. Saveski and A. Mantrach. Item cold-start recommendations: learning local collective embeddings. In *RecSys*, pages 89–96. ACM, 2014.

[28] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock. Methods and metrics for cold-start recommendations. In *SIGIR*, pages 253–260. ACM, 2002.

[29] M. Sharma, J. Zhou, J. Hu, and G. Karypis. Feature-based factorized bilinear similarity model for cold-start top-n item recommendation. In *SDM*, 2015.

[30] Y. Shi, M. Larson, and A. Hanjalic. List-wise learning to rank with matrix factorization for collaborative filtering. In *ACM RecSys*, pages 269–272. ACM, 2010.

[31] Y. Shi, M. Larson, and A. Hanjalic. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)*, 47(1):3, 2014.

[32] V. Sindhwani, S. S. Bucak, J. Hu, and A. Mojsilovic. One-class matrix completion with low-density factorizations. In *ICDM*, pages 1055–1060. IEEE, 2010.

[33] N. Srebro, J. Rennie, and T. S. Jaakkola. Maximum-margin matrix factorization. In *Advances in neural information processing systems*, pages 1329–1336, 2004.

[34] H. Steck. Training and testing of recommender systems on data missing not at random. In *KDD*, pages 713–722. ACM, 2010.

[35] H. Steck. Gaussian ranking by matrix factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 115–122. ACM, 2015.

[36] M. Volkovs and R. S. Zemel. Collaborative ranking with 17 parameters. In *Advances in Neural Information Processing Systems*, pages 2294–2302, 2012.

[37] M. Weimer, A. Karatzoglou, Q. V. Le, and A. Smola. Maximum margin matrix factorization for collaborative ranking. *NIPS*, 2007.

[38] T. Zhou, H. Shan, A. Banerjee, and G. Sapiro. Kernelized probabilistic matrix factorization: Exploiting graphs and side information. In *SDM*, volume 12, pages 403–414. SIAM, 2012.